27.1 A 500MHz Random Cycle 1.5ns-Latency, SOI Embedded DRAM Macro Featuring a 3T Micro Sense Amplifier

John Barth[1], William Reohr[3], Paul Parries[2], Greg Fredeman[2], John Golz[2], Stanley Schuster[3], Richard Matick[3], Hillery Hunter[3], Charles Tanner III[1], Joseph Harig[1], Hoki Kim[2], Babar Khan[2], John Griesemer[2], Robert Havreluk[3], Kenji Yanagisawa[2], Toshiaki Kirihata[2], Subramanian Iyer[2]

[1]IBM, Burlington, VT, [2]IBM, Hopewell Junction, NY
[3]IBM T.J Watson, Yorktown Heights, NY

Logic-based embedded DRAM has matured into a wide range of ASIC applications, SRAM replacements [1] and off-chip caches for microprocessors [2]. However, in on-chip microprocessor applications, embedded DRAM has been limited to moderate performance bulk logic technologies, and while it has been used in supercomputers such as BlueGene/L [3], it has not broken into the high-performance single-chip multi-core microprocessor arena [4]. This will happen only if embedded DRAM is developed on the same high-performance silicon-on-insulator (SOI) process platform and embedded DRAM random access and cycle times are dramatically improved. The availability of a high-performance embedded DRAM macro would allow cache-dominated chips [5] to double their cache capacity, while fitting in a smaller footprint. This enables significant system performance improvement[6]—simultaneously reducing cost, power consumption and soft-error vulnerability. This paper presents a prototype SOI embedded DRAM macro developed for high-performance microprocessors, and introduces a performance-enhancing, three-transistor micro-sense-amplifier (μSA) architecture. The macro is characterized via a testchip fabricated in a 65nm SOI DRAM process [7] that features a reduced-process-complexity adder compared with bulk embedded DRAM processes.

The 2Mb macro, shown in Fig. 27.1.1a, is composed of eight 292kb sub-arrays (256 WL×1216 BL) surrounding a central I/O region, where 10 addresses and 146 data are received/driven. Addresses AADD<0:1> select 1 of 4 sub-arrays on top and bottom in parallel. Word addresses WADD<0:1> select 1 of 4 master WLs, while WADD<2:7> select 1 of 64 global WLs. All are delivered to the sub-arrays on M4 and combined locally to select 1 of 256 local WLs. One-hot, late-select column signals support 16-way cache associativity. Eight are driven to the upper array, eight to the lower array, selecting 1 of 16 columns.

Figure 27.1.1b shows the unconventional orthogonal location of the WL driver as introduced in [8]. M4 is used in a pyramid-like pattern to *jumper* the orthogonal WL driver to the east/west M3 WL strap. The pyramid pattern, seen in the micrograph of Fig. 27.1.1c, is chosen to minimize WL skew by using the shortest M4 jumper on the longest M3 runs. The orthogonal WL configuration saves area by decoupling the WL driver layout from the array WL pitch and avoids replicating WL pre-decode in each sub- array. The I/O region design was reused from an existing SRAM design and therefore is not area-optimized.

Conventional techniques for improving DRAM performance involve reducing the BL length; however, this incurs area penalties from sense amplifiers, BL twisting and reference circuits. Sense-amplifier overhead is further degraded with body-tied SOI devices, required to prevent history-induced sense-amplifier mismatch.

Figure 27.1.2 introduces the μSA architecture, developed to provide high-performance sensing without incurring the overhead associated with conventional techniques. This hierarchical scheme relies on a high transfer ratio during a read to create a large voltage swing on a local bit-line (LBL)—large enough to be sampled with the single-ended μSA. Each μSA services 32 cells connected on a single, relaxed-pitch M1. The μSA transfers data to/from a secondary sense amplifier (SSA) via two global read/write BLs (RBL/WBL). Each SSA, in turn, services 8 μSAs and transfers data to/from a tertiary sense amp (TSA) via local M1 data lines, again in a hierarchical fashion, as shown in Fig. 27.1.3.

To achieve a high transfer ratio, a 20fF deep-trench cell is used in combination with a 4fF LBL. The low capacitance LBL is achieved using a short 32b LBL. The single-ended LBL enables relaxed M1 pitch, increasing line-to-line space by 3×. Low junction capacitance, inherent with SOI, further reduces LBL capacitance. The low LBL capacitance results in a reduced read RC time constant, simultaneously improving access and cycle times. The resulting capacitance ratio allows 83% of the cell's voltage to be transferred to the LBL, ideally providing 830mV to the μSA.

The μSA in Fig. 27.1.2 is configured with an NMOS for reading, an NMOS for driving the LBL to a full low level and a PMOS for driving the LBL to a full high level. RBL/WBL, routed on minimum M2, parallel to M1 LBLs, control read/write operations to all μSAs in parallel. All cycles start in the pre-charge condition with equalization signals (BEQN/ SEQN) low, holding RBL and WBL high, precharging LBLs to ground. Sub-array selection first triggers SEQN high, forcing WBL low, floating LBL and enabling μSA NMOS read heads. BEQN is held low until WBL has fallen low, absorbing any line-to-line coupling from WBLs to RBLs. After WBL falls, BEQN rises, floating RBL and enabling read or write operations.

Write data is delivered to the tertiary sense amp (TSA), shown in Fig. 27.1.3, via M4 global data lines (GDLT/GDLC). Initially precharged low, GDLT or GDLC is remotely driven high during a write. To write a '1', GDLT is driven high, TSA drives local data line LDLC low, passes to the column selected by CSL<0:7> and pulls RBL low. This forces LBL high, writing a '1' into the node of the selected cell. To write a '0', GDLC is remotely driven high, TSA drives LDLT low and passes to LT of the SSA selected by CSL<0:7>. This forces WBL and RBL high, driving LBL low without contention, writing a '0' into the selected cell. Fig. 27.1.4 shows RBL/WBL/LBL/Node/WL waveforms of write and read cycles.

Read data is transferred from the cell to the LBL upon activation of the selected WL. For a stored '1', LBL rises at least 1 threshold above the read transistor of the associated μSA, weakly pulling RBL low. When RBL falls below the threshold of the μSA PMOS, feedback drives the LBL to a full high level. This amplifies LBL charge, refreshes the cell, and strongly drives RBL low. The refresh of a '1' is self-timed, requiring no external control. RBL falling passes from the selected column to the TSA, driving GDLT high. For a stored '0', LBL remains low, RBL remains high and WBL remains low until external timing signal SETP triggers SSA to evaluate RBL. With RBL high, LT falls, driving WBL high, driving LBL low, refreshing the '0'. LT's fall simultaneously flows through the SSA, from the selected column to the TSA, driving GDLC high.

TSA, SSA and μSA are configured to be bi-directional and support direct write [9] of data prior to WL activation, without disrupting reads of adjacent sense amplifierss and data bits. Even with fast, self-timed refresh of a '1', write '1' remains the cycle limiter. Consequently, direct write is required to achieve minimum cycle time. In contrast to SRAM cells and DRAM sense amplifiers, there is no reliance on NMOS/PMOS strength ratios for writing to the opposite state. This enables very low 600mV operation, limited only by the highest threshold device.

The shmoo plot of Fig. 27.1.5 demonstrates 1.5ns access time with 99.99% bit yields at 1V 85°C across the full address space. The design achieves 4ns access at 600mV. Figure 27.1.6 demonstrates 500MHz operation at 1V over a 32k address space (limited by high frequency tester). Figure 27.1.7 includes a table of features and the chip micrograph of 2Mb and 4Mb macros.

References

[1] H. Pilo, D. Anand, J. Barth, et al, "A 5.6-ns Random Cycle 144-Mb DRAM with 1.4Gb/s/pin and DDR3-SRAM Interface", *IEEE J. Solid-State Circuits*, vol. 38, no. 11, pp. 1974-1980, Nov. 2003,.

[2] J. Clabes, J. Friedrich, M. Sweet, et al., "Design and Implementation of the Power5 Micro Processor," *ISSCC Dig. Tech. Papers*, pp. 56-57, Feb., 2004.

[3] S. Iyer, J. Barth, P. Parries, et al., "Embedded DRAM: Technology Platform for Blue Gene/L Chip", *IBM J. Research Dev.,* vol. 49, no. 2/3, , pp. 333-350, Mar./May, 2005.

[4] E.B. Cohen, N.J. Rohrer, P. Sandon, et al., "A 64B CPU Pair: Dual and Single-Processor Chips," *ISSCC Dig. Tech. Papers*, pp. 106-107, Feb., 2005.

[5] S. Naffziger, B. Stackhouse, and T. Grutkowski, "The Implementation of a 2-Core Multi-Threaded Itanium-Family Processor, *ISSCC Dig. Tech. Papers*, pp. 82-83, Feb., 2005.

[6] R. Matick, S. Schuster, "Logic-based eDRAM: Origins and Rationale for Use", *IBM J. Research Dev.,* vol. 49, no. 1, pp. 145-165, Jan., 2005.

[7] G. Wang, et al., "A 0.127 μm² High Performance 65nm SOI Based Embedded DRAM for On-Processor Applications," *IEDM,* Dec., 2006, accepted for publication.

[8] D. Weiss, J.J. Wuu, and V. Chin, "The On-chip 3-MB Subarray-Based Third-Level Cache on an Itanium Microprocessor," *ISSCC Dig. Tech. Papers*, pp. 112-113, 2002.

[9] J. Barth, D. Anand, J. Dreibelbis, et al., "A 500 MHz Multi-Banked Compilable DRAM Macro with Direct Write and Programmable Pipeline," *ISSCC Dig. Tech. Papers*, pp. 204-205, Feb., 2004.

[10] T. Kirihata, P. Parries, D. Hanson, et al., "An 800MHz Embedded DRAM with a Concurrent Refresh Mode," *ISSCC Dig. Tech. Papers*, pp. 206-207, Feb., 2004.

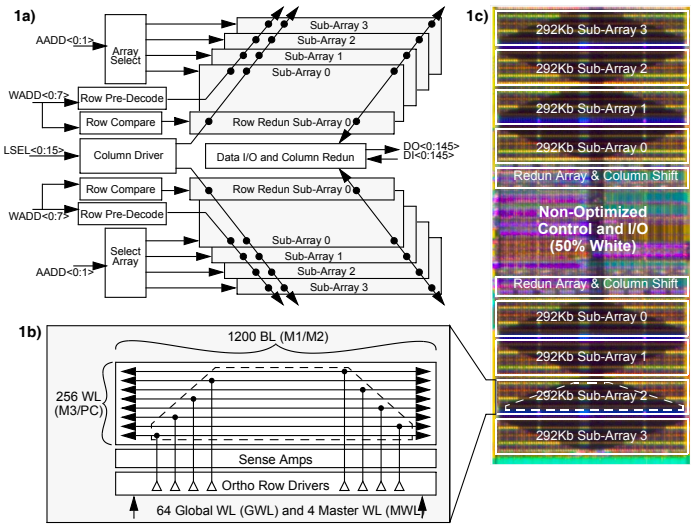

Figure 27.1.1: Macro Architecture.

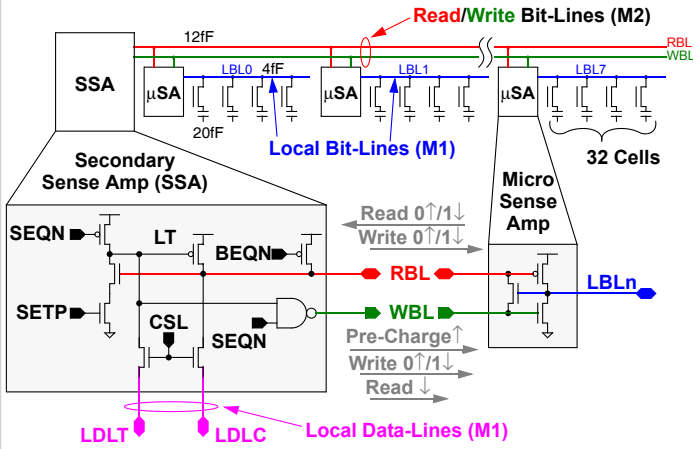

Figure 27.1.2: Micro Sense Architecture (µSA).

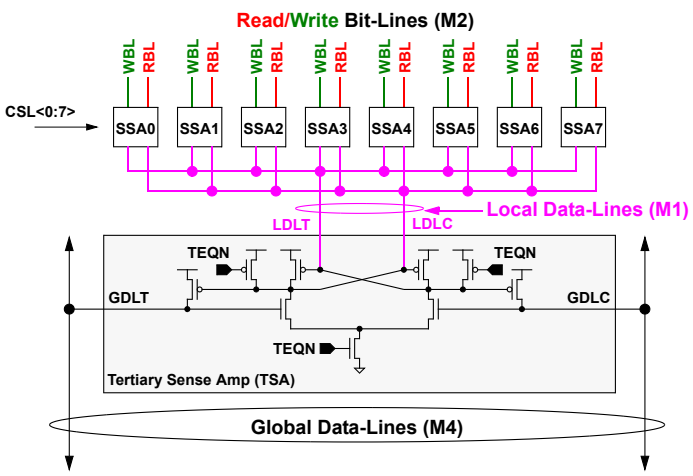

Figure 27.1.3: Tertiary Sense Amp.

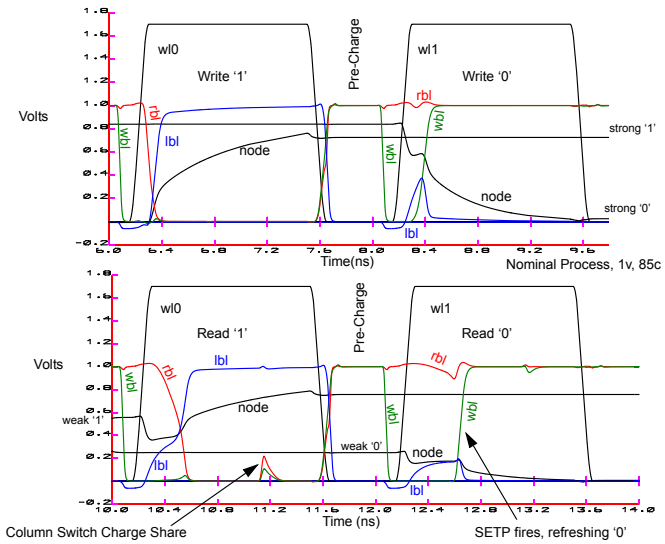

Figure 27.1.4: Write/Read Simulation.

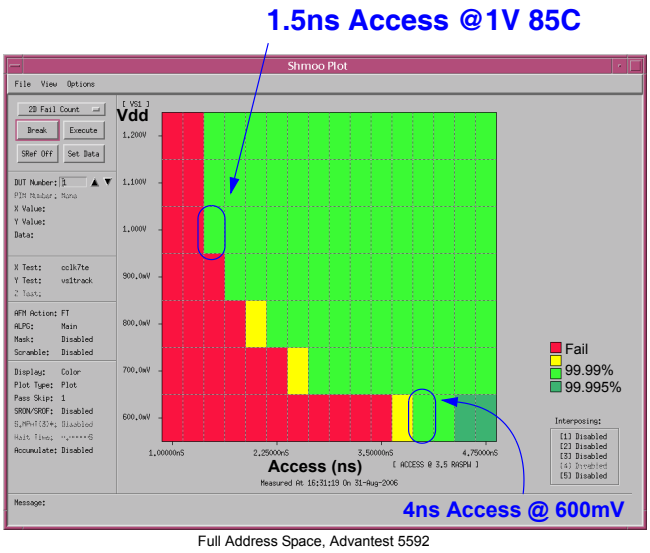

Full Address Space, Advantest 5592

Figure 27.1.5: Access Shmoo.

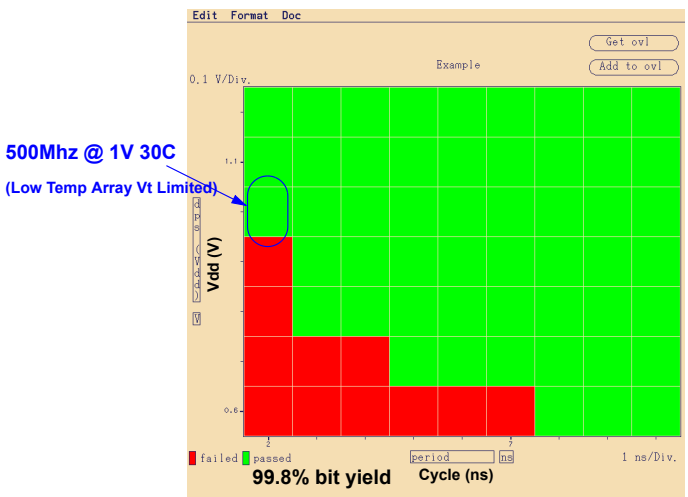

Figure 27.1.6: Random Cycle Shmoo.

Continued on Page 617

Technology	65nm PD-SOI (Tox=1.12nm/2.35nm)
Cell Size	0.23μmx0.55μm(0.1265μm^2)
Retention	40us @ 105c, 98.7% Availability†
Macro Size	568μmx1170μm†† 0.665mm^2 (0.317x SRAM)
Organization	1K Row x 16 Col x 146 (2Mb)
Sub-Array	Hierarchical w/ Micro Sense 32 Rows/Local BL 8 Local BL/Global BL 1200 Bits/WL
Redundancy	2 - 32 WL Row Arrays Static Column Shifting
Supply	1.0V / 1.7V (WL High)
Performance	2ns Random Cycle 1.5ns Random Access
Power†††	AC only - 76mW (0.84x SRAM) Stdby + Refresh - 42mW (0.18x SRAM)
Notes	† With Concurrent Refresh [10] †† White Space removed from Control ††† Calculated, Nominal Process, 1.0V

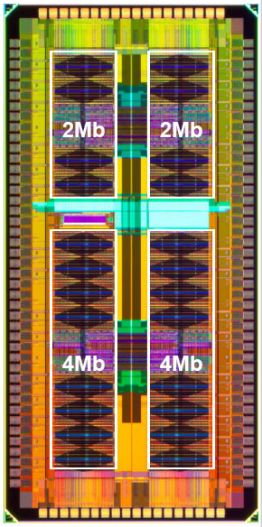

Figure 27.1.7: Chip Micrograph and Features.

